

Weekly Report Week 1 & 2

Wen Han Chia

Lizard Classification Project

January 17, 2025

1 Time Log

1.1 What did you do this week?

1. Registered for CS8903 Special Problems Module
2. Organized bi-weekly meetings with computational advisor
3. Reviewed Dr Stroud' s meeting with computational advisors
4. Conducted Literature Review on classification model and solutions to tackle imbalanced dataset
5. Drafted project scope on the Methods document based on the information above
6. Accessed and forked GitHub repository of the previous Lizard Classification Project

1.2 What are you going to do next week?

- Download Anole Species dataset from iNaturalist
- Iterate on Methods document if needed
- Continue conducting literature review on classification model and class imbalance

2 Abstracts

Y. Cui, M. Jia, T. -Y. Lin, Y. Song and S. Belongie, "Class-Balanced Loss Based on Effective Number of Samples," *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA, 2019, pp. 9260-9269, doi: [10.1109/CVPR.2019.00949](https://doi.org/10.1109/CVPR.2019.00949)

With the rapid increase of large-scale, real-world datasets, it becomes critical to address the problem of long-tailed data distribution (i.e., a few classes account for most of the data, while most classes are under-represented). Existing solutions typically adopt class re-balancing strategies such as re-sampling and re-weighting based on the number of observations for each class. In this work, we argue that as the number of samples increases, the additional benefit of a newly added data point will diminish. We introduce a novel theoretical framework to measure data overlap by associating with each sample a small neighboring region rather than a single point. The effective number of samples is defined as the volume of samples and can be calculated by a simple formula $(1-\beta n)/(1-\beta)$, where n is the number of samples and $\beta \in [0, 1)$ is a hyperparameter. We design a re-weighting scheme that uses the effective number of samples for each class to re-balance the loss, thereby yielding a class-balanced loss. Comprehensive experiments are conducted on artificially induced long-tailed CIFAR datasets and large-scale datasets including ImageNet and iNaturalist. Our results show that when trained with the proposed class-balanced loss, the network is able to achieve significant performance gains on long-tailed datasets

3 What you did and proof

These 2 weeks mostly consisted of onboarding into the team, familiarizing with the project and finding out what had been previously done, and drafting the project's scope. I am still in the midst of reviewing the codebase and understanding the work previously done. Below is a screenshot of the codebase and a screenshot of the draft of the Methods submission.

```
9
10 """
11 Training script for the 5 class classifier.
12 """
13
14 # Load object detection model
15 detection_model = tf.keras.models.load_model('F:/LizardCV/detection.h5')
16
17 # Load the MobileNetV2 model pre-trained on ImageNet
18 base_model = MobileNetV2(input_shape=(224, 224, 3), include_top=False, weights='imagenet')
19
20 # Freeze the base model (i.e., make its weights non-trainable)
21 base_model.trainable = False
22
23 # Add custom classification head
24 model = models.Sequential([
25     base_model,
26     layers.GlobalAveragePooling2D(),
27     layers.Dense(128, activation='relu'),
28     layers.Dropout(0.5),
29     layers.Dense(5, activation='softmax')
30 ])
31
32 # Compile the model
33 model.compile(optimizer=Adam(learning_rate=0.00001),
34              loss=tf.keras.losses.CategoricalCrossentropy(label_smoothing=0.1),
35              metrics=['accuracy', tf.keras.metrics.Recall(), tf.keras.metrics.Precision()])
36
37 # Data loading
38 dataset, class_weight_dict = load_dataset_with_labels('F:/LizardCV/Raw', None, 10000)
39 train_dataset = dataset.map(lambda image, label, id: (tf.image.resize(image, [224, 224]), label))
40 #train_dataset = train_dataset.map(lambda image, label: (
41 #    tf.image.resize(tf.image.random_flip_left_right(tf.image.random_brightness(image, 0.2)), [224, 224]), label
42 #))
43 # Define EarlyStopping callback to monitor validation loss
44 early_stopping_callback = tf.keras.callbacks.EarlyStopping(
45     monitor='loss', # You can also use 'val_accuracy' or another metric
46     patience=5, # Number of epochs to wait for improvement
47     restore_best_weights=True
48 )
49
50 # Train the model
51 history = model.fit(
52     train_dataset,
53     epochs=20,
54     class_weight=class_weight_dict,
55     callbacks=[early_stopping_callback]
56 )
```

Figure 1: A screenshot of the codebase of the previous student's work

METHODS			
PROJECT OUTLINE			
<p>For each applicable step in the project cycle provide the (1) actions you will take (2) the justification for that action (3) the action you will take if the initial action fails. We are looking only for action to be included in your final publication.</p> <p>Example Action: Annotate 1000 frames with bounding boxes for male and female fish</p> <p>Not an Action: Read literature about computer vision and fish.</p>			
Data Processing			
Action	Justification	Plan if Action Fails	Status
Ensure proper labelling of all images from datasets	This will be very helpful in training image classification models	Seek guidance from students/advisors on the method used previously to retrieve dataset	Not started
Over-sample from minority classes and under-sample from majority classes	This increases the minority class count and reduces the majority class counts to reduce the effect of class imbalance.	Use data augmentation to increase count of minority classes by doing minor adjustments to images (cropping, rotation, scaling)	Not started
Algorithm Development			
Action	Justification	Plan if Action Fails	Status
Assign higher weights to minority classes in the loss function	This increases the weight of the minority classes, thus changing the loss function output w.r.t to these classes	Revert to default loss function	
Using different loss function (e.g. Focal Loss)	Shown to address class imbalance issue by focusing on hard-to-classify samples		
Hyperparameter-tuning: train models with different hyper parameters and determine ideal set of hyperparameters (batch size, learning rate, optimizer)			
Model Selection/ Model Training			
Action	Justification	Plan if Action Fails	Status
Train a classification model from scratch (e.g. ResNet, Vision Transformers) to classify Anole species	Training the entire model from scratch (instead of fine-tuning) allows the model's weights to be tuned to object of interest (anole species). This could potentially lead to better accuracy	Combine plans from row 1 and 2 - Fine-tune a detection model, and train the classification model using cropped image generated from detection model. Finally, combine these 2 models into a single pipeline.	Not started
Develop a pipeline consisting of object detection, extracting the region of interest, and passing cropped image to classification models	Object detection model provides location of object of interest before cropping the bounding box. Image processing techniques (sharpen) can be performed on the cropped region of interest before passing to classification model. This could result in more accurate classification	Refer to last row	Not started
Validation of method			
Action	Justification	Plan if Action Fails	Status
For object detection model, we will use precision, recall, mean Average Precision (mAP), and Intersection over Union (IoU) to evaluate model's performance	mAP and IoU evaluates the model's accuracy in localizing area of interest		Not started
For classification models, we will use F1 score and Precision-Recall curves to evaluate model's performance	F1 score and precision-recall curves are preferred metrics when dealing with imbalanced datasets		Not started
User-interface/ User-experience development			
Action	Justification	Plan if Action Fails	Status
A basic website that enables user to upload a photo and classify the species.	Ensuring classification model has reasonable accuracy will be prioritized first		
Fine-tune end-to-end detection and classification models (e.g. Yolov8 and Faster R-CNN)	These models are capable of performing object localization and classification. Thus, we only need to fine-tune these model to detect and classify specific anole species	Refer to other plans listed in other rows	

Figure 2: A screen shot of the Methods submission draft